Contemporary Epistemic Logic and the Lockean Thesis

Lorenz Demey

1 Introduction

Classical epistemology mainly uses *qualitative* notions, such as knowledge, belief, justification, etc. (Williams, 2001). Formal epistemology, however, makes extensive use of *quantitative* notions, such as degrees of belief, coherence, confirmation, etc. (Douven and Meijs, 2007; Eels and Fitelson, 2000; Huber and Schmidt-Petri, 2009). A natural question thus arises: what is the relationship—if any—between the qualitative and the quantitative framework? In particular, one may ask whether there is a relation between *belief* and *degrees of belief*.

A widespread thesis about this issue is that the qualitative notion of belief is reducible to the quantitative notion of degree of belief: believing that φ is defined as having a 'sufficiently high' degree of belief that φ. Foley (1992) has labeled this the 'Lockean thesis'. The main aim of this paper is to explore the advantages and disadvantages of this thesis from the perspective of contemporary epistemic logic. I will argue that, although the Lockean thesis is quite problematic for *classical* epistemic logic, it seems to have a much brighter future in *contemporary* epistemic logic.

To keep the paper relatively self-contained, Sections 2 and 3 provide brief overviews of the two central themes. In Section 2 I discuss the specific features of *contemporary* epistemic logic, and compare them with those of *classical* epistemic logic. Section 3 introduces the formal details of the Lockean thesis. This thesis yields a notion of belief which is not closed under conjunction. I will discuss the relationship of this problem with the well-known lottery paradox. Section 4 is the core section of this paper. I argue that the conjunction problem is typical for *classical* epistemic logic, and propose to reconsider this thesis from the perspective of *contemporary* epistemic logic. In particular, I will focus on the dynamic behavior of the notion of belief obtained via the thesis. To this end, I will introduce a system of public announcement logic, enriched with a (qualitative) belief operator, and a system of probabilistic public announcement logic (in which the Lockean thesis can be applied to 'define' a belief operator). It turns out that accepting the Lockean thesis leads to a unified perspective on the dynamic behavior of belief and degrees of belief, which illustrates its methodological fruitfulness. Furthermore, I will argue that, when combined with Baltag's so-called 'Erlangen program' in epistemology, this observation also constitutes a philosophical argument in favor of the Lockean thesis. Section 5, finally, summarizes the results obtained in this paper, and suggests some questions for further inquiry.

2 Contemporary Epistemic Logic

The aim of this section is to introduce and discuss the most important features of contemporary epistemic logic, and to compare them with those of classical epistemic logic.[1] First, however, it should be emphasized that, despite the terminology ('classical'/'contemporary') being used, the distinction being made is

[1]This section presupposes a basic familiarity with the central topics in epistemic logic. I will mention several formal notions and theorems, without going into any detail: they merely serve to illustrate the large-scale distinction introduced in this section (classical/contemporary epistemic logic). Technical details will, where necessary, be introduced in Section 4.

in the first place a *conceptual* one, rather than a strictly *historical* one. Most work on classical epistemic logic was being done before the emergence of contemporary epistemic logic, but one can certainly find examples of contemporary epistemic logic as early as the late 1960's (Lewis, 1969), and conversely, some logicians are still doing (very valuable) work in classical epistemic logic today (Halpern et al., 2009).

The starting point of classical epistemic logic, and of epistemic logic in general, is Hintikka's seminal *Knowledge and Belief* (1962). In this work, knowledge is analyzed as a modal operator, which is given a semantics in terms of Kripke models. The formula $K_i\varphi$ thus means: 'agent i knows that φ'. Hintikka used his framework to gain insight about principles such as $K_i\varphi \rightarrow K_iK_i\varphi$ (if agent i knows that φ, does it then follow that she knows that she knows this?). Two features are of central importance in this framework.

First, the framework is essentially *single-agent*. It is about the knowledge of one single agent, not about the (pieces of) knowledge of several agents, and how these might interact. One can trivially go from one to many agents, by simply 'adding subscripts'; for example, one then gets formulas such as $K_i\varphi \wedge \neg K_j\varphi$ ('agent i knows that φ, but agent j doesn't'). However, in this way one still cannot obtain the irreducibly social notions of *common knowledge* and *distributed knowledge*, i.e. these notions cannot be defined in terms of the individual knowledge operators (Halpern and Moses, 1990).

Second, the framework is *static*. It focuses entirely on an agent's knowledge at a single point in time, without taking into consideration that the agent's knowledge might change over time (e.g. because she learns about new information). For example, Hintikka explicitly rules out occasions "on which people are engaged in gathering new factual information. Uttered on such an occasion, the sentences 'I don't know whether p' and [later] 'I know that p' need not be inconsistent" (1962, p. 7–8).

Contemporary epistemic logic (as this term is used here) can be defined as the opposite of classical epistemic logic with respect to exactly these two key features.

In the first place, contemporary epistemic logic is a *multi-agent* enterprise. Because of applications in economics and computer science (distributed systems), the notion of common knowledge has become very important, and several characterizations of this notion are available (the most important ones being the iterative and the fixed-point characterization) (Barwise, 1988). Similarly, the notion of distributed knowledge has been studied extensively (van der Hoek et al., 1999).

In the second place, contemporary epistemic logic focuses on the *dynamics* of knowledge. One typically studies scenarios that involve learning: at first, an agent does not know whether φ; next, φ is (truthfully) announced; then, after the announcement, the agent *does* know that φ. Dynamic epistemic logic can be used to formalize and analyze such scenarios, but also more complicated ones, such as card games. It is therefore often applied in computer science (protocol security) (van Ditmarsch et al., 2007).

Finally, it should be noted that these two themes (multi-agent/dynamics) often interact with each other. For example, the distinction between three important types of epistemic dynamics, viz. public announcements, private announcements, and semi-private announcements (Baltag and Moss, 2004), only makes sense in a multi-agent setting: in a single-agent setting these three notions collapse into each other.[2]

[2]It is interesting to note that this double evolution (from single-agent to multi-agent, from static to dynamic) has taken place not only in epistemic logic, but also in epistemology. Classical epistemology deals with the question whether a single agent, at a given point in time, does or does not possess knowledge concerning some proposition φ. Contemporary epistemologists, however, also deal with *social* (multi-agent) phenomena such as knowledge by testimony and the role of experts (Goldman, 1999); furthermore, formal epistemologists study how new infor-

3 The Lockean Thesis

Classical epistemology mainly uses the qualitative notion of *belief*, whereas formal epistemology makes extensive use of the quantitative notion of *degrees of belief*. In most work, degrees of beliefs are formalized as (subjective) probabilities.[3] One thus works with statements such as $P(\varphi) = k$ (for $k \in [0, 1]$), which means that the agent assigns probability k to proposition φ.

A widespread thesis, sometimes called the *Lockean thesis*, is that "it is epistemically rational for us to believe a proposition just in case it is epistemically rational for us to have *sufficiently high* degree of confidence in it" (Foley, 1992, p. 111, my emphasis). Formally, this means that in a purely probabilistic framework, one can define ('qualitative') belief as follows:

$$B\varphi :\equiv P(\varphi) \geq \tau. \tag{1}$$

Here, τ is a *treshold*: a degree of belief is 'sufficiently' high to count as a (qualitative) belief iff that degree of belief is above τ.

There has been a lot of discussion about what the exact value of the treshold τ should be. There seems to be a consensus that τ should be at least 0.5. If $\tau = 0.4$, for example, then for a proposition φ with $P(\varphi) = 0.45 \geq \tau$, thesis (1) yields $B\varphi$, but it also follows that $P(\neg\varphi) = 1 - 0.45 = 0.55 \geq \tau$, and thus also $B\neg\varphi$, i.e. the resulting notion of belief would allow for inconsistent beliefs. If $\tau > 0.5$ this case cannot occur: if $B\varphi$, then $P(\varphi) \geq \tau > 0.5$, and thus $P(\neg\varphi) = 1 - P(\varphi) < 0.5 < \tau$, i.e. $\neg B\neg\varphi$. This principle (if $B\varphi$ then $\neg B\neg\varphi$) is sometimes called the *consistency requirement* about belief.

mation should be processed (e.g. via Bayesian updating or Jeffrey conditionalization (Jeffrey, 1983)). The methodological consequences of this analogy will be explored in future work (it seems to suggest a unified perspective on epistemic logic and epistemology).

[3]However, degrees of belief can also be formalized in non-probabilistic frameworks, such as possibility theory and ranking theory (Dubois and Prade, 2009; Spohn, 2009).

Some authors have proposed to take $\tau = 1$, but this seems to be too strong: belief intuitively does not require complete certainty. For $\tau < 1$, however, a well-known problem for the Lockean thesis arises, viz. the resulting notion of belief is not closed under conjunction. For example, suppose that $\tau = 0.6$, consider a fair six-faced die, write p for 'the die will land with 1,2,3 or 4 eyes up' and q for 'the die will land with 3,4,5 or 6 eyes up'; then $P(p) = P(q) = 0.66 \geq \tau$ and $P(p \wedge q) = 0.33 < \tau$, and thus (1) yields Bp and Bq, while $\neg B(p \wedge q)$. One might think that this problem can be solved by taking τ to be increasingly closer to 1, e.g. 0.95. However, consider a fair lottery with 100 tickets (the agent considers all tickets equally likely to win, and exactly one ticket will win) and write p_i for 'ticket i will not win'; then $P(p_i) = 0.99 \geq 0.95$ for each i, while $P(\bigwedge_{i=1}^{i=100} p_i) = 0 < 0.95$, and thus (1) yields Bp_i for each i, while $\neg B(\bigwedge_{i=1}^{i=100} p_i)$.

This explains the central role of the lottery paradox in this context: fair lotteries form a canonical class of counterexamples to belief being closed under conjunction. No matter how close to 1 the value of τ is taken to be, one can construct a fair lottery (with a sufficiently large number of tickets) which yields a finite number of propositions that are believed (i.e. their probabilities are $\geq \tau$), while their conjunction is not believed (i.e. its probability is $< \tau$).

4 A Dynamic Look at the Lockean Thesis

In the previous section I introduced the Lockean thesis, and discussed its most important problem, viz. that it yields a notion of belief which is not closed under conjunction. It is clear that this problem typically belongs to *classical* epistemic logic (cf. Section 2): it is about a single agent, and it is a static scenario (the agent's beliefs are examined at a single point in time).

One can also ask, however, how the Lockean thesis fares from the perspective of *contemporary* epistemic logic. This means that one should ask questions such as: does this thesis give rise to interesting social (*multi-agent*) notions of belief, e.g. common belief? In the game-theoretical literature, the notion of belief obtained via the Lockean thesis is often called p-belief (because game theorists usually use the letter p, rather than τ, to denote the treshold). Just like (qualitative) belief (and knowledge) can be used to define common belief (and common knowledge), the notion of p-belief gives rise to a notion of *common p-belief*. The formal behavior of common p-belief largely resembles that of 'classical' (qualitative) common belief; for example, it has both an iterative and a fixed-point characterization (Kajii and Morris, 1997; Monderer and Samet, 1989). Furthermore, many 'applications' that require the notion of common belief can equally well be modeled using the notion of common p-belief; a typical example is the agreeing to disagree theorem in game theory.[4] To summarize: both from a theoretical and an application-oriented[5] perspective, the Lockean thesis seems to transfer well from the single-agent to the multi-agent case.

In this section, however, I will focus on the other characteristic feature of contemporary epistemic logic: *dynamics*. I will study the dynamic behavior of the notion of belief generated by the Lockean thesis, and compare it with the dynamic behavior of a 'classical' qualitative notion of belief. The focus will be on one particular type of dynamics, viz. public announcements.

Subsection 4.1 introduces a system of public announcement logic, enriched with a (qualitative) belief operator. Subsection 4.2 introduces probabilistic public announcement logic, in which the Lockean thesis can be applied to 'define' a belief operator. In Subsection 4.3 I will compare the dynamic behavior of the qualitative belief operator and the probabilistically defined belief operator,

[4] The first version of the agreeing to disagree theorem involved common *knowledge* and was proved by Aumann (1976). Dégremont and Roy (2009) established a version of this theorem that only requires the 'classical' (qualitative) notion of common *belief*. Monderer and Samet (1989) established a version of this theorem using the notion of common *p-belief*.

[5] I will return to this application-oriented perspective on the Lockean thesis; cf. Footnote 12.

and argue that the results of this comparison constitute a methodological and a philosophical argument in favor of the Lockean thesis.

4.1 Public Announcement Logic with Beliefs

I will now give a brief overview of a system of public announcement logic, enriched with a belief operator. It is well-known that such systems cannot plausibly be interpreted on Kripke models: if an agent receives a true piece of information φ while previously believing that $\neg\varphi$, then this agent is predicted to go insane and start believing *everything* (rather than performing a realistic process of *belief revision*)—thus contradicting the consistency requirement about belief.[6] Therefore, systems of public announcement logic with a belief operator are interpreted on epistemic plausibility models (Baltag and Smets, 2008; van Benthem, 2007).

First I introduce the static part. Fix a finite set I of agents, and a denumerably infinite set $Prop$ of proposition letters. An *epistemic plausibility model* (in the sense of Baltag and Smets)[7] is defined to be a structure $\mathbb{M} := \langle W, \sim_i, \leq_i, V\rangle_{i\in I}$, where W is a nonempty set of states, $V : Prop \to \wp(W)$ is a valuation, for each $i \in I$, $\sim_i$ is an equivalence relation on W (which gives rise to equivalence classes $[w]_{\sim_i} := \{v \in W \mid w \sim_i v\}$), and $\leq_i$ is a relation on W satisfying the following two conditions:

1. if $s \leq_i t$ then $s \sim_i t$,
2. the restriction of $\leq_i$ with respect to each $\sim_i$-equivalence class (formally: $\leq_i \cap ([w]_{\sim_i} \times [w]_{\sim_i})$, for each $w \in W$) is a well-preorder.

[6] Cf. van Benthem (2007, Section 3.1) for more details.

[7] Baltag and Smets (2008) and van Benthem (2007) use subtly different definitions of epistemic plausibility model. In (Demey, 2011b) I compare the model theory of both notions and propose a methodological argument in favor of Baltag and Smets's notion.

For any state $w \in W$ and set $X \subseteq [w]_{\sim_i}$, the set of $\leq_i$-minimal elements of X is defined as $\mathrm{Min}_{\leq_i}(X) := \{x \in X \mid \forall y \in X : x \leq_i y\}$. That $\leq_i \cap ([w]_{\sim_i} \times [w]_{\sim_i})$ is a well-preorder means that it is reflexive and transitive, and that for each nonempty $X \subseteq [w]_{\sim_i}$ also $\mathrm{Min}_{\leq_i}(X)$ is nonempty.

The language that will be interpreted on these structures is a propositional language with a knowledge- and a belief-operator. However, for technical reasons (which will be discussed later), this language does not contain an *ordinary belief* operator, but rather a *conditional belief* operator $B_i(\cdot\,|\,\cdot)$. A formula such as $B_i(\varphi|\psi)$ should be read as: 'agent i believes that φ, conditional on ψ'. The ordinary belief operator is definable in terms of the conditional belief operator, by putting $B_i\varphi :\equiv B_i(\varphi|\top)$ (where $\top$ denotes an arbitrary tautology).

I now turn to the semantics of this language. The formal semantic clauses for the proposition letters and the Boolean connectives are standard; the formal semantic clauses for the knowledge operator and the the conditional belief operator are the following ($[\![\,\varphi\,]\!]^{\mathbb{M}}$ abbreviates the set $\{w \in W \mid \mathbb{M}, w \models \varphi\}$):

$$\begin{array}{lll} \mathbb{M}, w \models K_i\varphi & \text{iff} & \forall v \in [w]_{\sim_i} : \mathbb{M}, v \models \varphi, \\ \mathbb{M}, w \models B_i(\varphi|\psi) & \text{iff} & \forall v \in \mathrm{Min}_{\leq_i}([w]_{\sim_i} \cap [\![\,\psi\,]\!]^{\mathbb{M}}) : \mathbb{M}, v \models \varphi. \end{array}$$

A sound and complete axiomatization of this logic is readily available; cf. Baltag and Smets (2008, Section 2.5).

I now introduce the dynamic part. Given an epistemic plausibility model $\mathbb{M} = \langle W, \sim_i, \leq_i, V\rangle_{i\in I}$ and a formula φ which is true in at least one state in $\mathbb{M}$, the *updated model* $\mathbb{M}|\varphi := \langle W^\varphi, \sim_i^\varphi, \leq_i^\varphi, V^\varphi\rangle$ is defined as follows:

- $W^\varphi = [\![\,\varphi\,]\!]^{\mathbb{M}}$,
- for every $i \in I$: $\sim_i^\varphi = \sim_i \cap (W^\varphi \times W^\varphi)$ and $\leq_i^\varphi = \leq_i \cap (W^\varphi \times W^\varphi)$,
- for every $p \in \mathit{Prop}$: $V^\varphi(p) = V(p) \cap W^\varphi$.

(It is straightforward to check that the updated model $\mathbb{M}|\varphi$ is indeed an epistemic plausibility model, as defined above.) The language is extended with a dynamic operator $[!\,\cdot]$. A formula such as $[!\varphi]\psi$ should be read as: 'after any public announcement of φ, it will be the case that ψ'. The dual of $[!\varphi]\psi$ is $\langle !\varphi\rangle\psi$, which is defined as $\neg[!\varphi]\neg\psi$, and which should be read as: 'φ can be announced, and afterwards it will be the case that ψ'. The formal semantic clauses look as follows:

$$\begin{array}{lll} \mathbb{M}, w \models [!\varphi]\psi & \text{iff} & \text{if } \mathbb{M}, w \models \varphi \text{ then } \mathbb{M}|\varphi, w \models \psi, \\ \mathbb{M}, w \models \langle !\varphi\rangle\psi & \text{iff} & \mathbb{M}, w \models \varphi \text{ and } \mathbb{M}|\varphi, w \models \psi. \end{array}$$

To obtain a sound and complete axiomatization for this dynamified logic, one merely needs to add 'reduction axioms'. These are biconditional statements which allow us to recursively rewrite formulas containing dynamic operators as formulas without such operators; hence the dynamic logic is equally expressive as the static one, and proving completeness for the dynamic logic can be reduced to that of the static logic (van Ditmarsch et al., 2007). Alternatively, reduction axioms can be seen as 'predicting' what will be the case *after* the dynamics has taken place in terms what is the case *before* the dynamics has taken place.

For expository purposes, I first state the reduction axiom for the ordinary belief operator:

$$[!\varphi]B_i\psi \longleftrightarrow \Big(\varphi \to B_i(\langle !\varphi\rangle\psi \mid \varphi)\Big). \tag{2}$$

This illustrates the two perspectives on reduction axioms discussed above. First of all, when (2) is read 'from left to right', it states that the public announcement operator $[!\varphi]$ can be 'pushed through' the complex formula $B_i\varphi$: on the right-hand side its scope is just ψ, which has a lower complexity than the original $B_i\psi$. Using the other reduction axioms as well, one can thus rewrite $[!\varphi]B_i\psi$ as a formula that does not involve the public announcement operator at all. Secondly, when (2) is read 'from right to left', it 'predicts' that agent i will believe that ψ *after* the public announcement of φ, just in case *before* the announcement, she believed $\langle !\varphi\rangle\psi$, conditional on φ.

Note that (2), which is the reduction axiom for the *ordinary* belief operator, requires the *conditional* belief operator; this is (one of) the reason(s) for introducing this conditional belief operator from the start (cf. supra). The reduction axiom for the conditional belief operator looks as follows:

$$[!\varphi]B_i(\psi|\alpha) \longleftrightarrow \Big(\varphi \to B_i(\langle !\varphi\rangle\psi \,|\, \langle !\varphi\rangle\alpha)\Big). \tag{3}$$

4.2 Probabilistic Dynamic Epistemic Logic

I will now give a brief overview of probabilistic public announcement logic (Demey, 2010; Kooi, 2003). Again, I begin with the static part. The sets I and *Prop* are used as in the previous subsection. A *probabilistic Kripke model* is defined to be a structure $\mathbb{M} := \langle W, \sim_i, \mu_i, V\rangle_{i\in I}$, where W is a nonempty, finite set of states, $V\colon Prop \to \wp(W)$ is a valuation, for each $i \in I$, $\sim_i$ is an equivalence relation on W, and $\mu_i\colon W \to (W \to [0,1])$ assigns to each state $w \in W$ a probability mass function $\mu_i(w)\colon W \to [0,1]$, satisfying the following conditions:[8]

1. $\mu_i(w)(w) > 0$ for all states $w \in W$,
2. $\mu_i(w)(v) = 0$ for all states $w, v \in W$ such that $w \not\sim_i v$.

On such structures, one can interpret formulas of the form $a_1P_i(\varphi_1) + \cdots + a_nP_i(\varphi_n) \geq k$, where $n \in \mathbb{N}$ and $a_1, \ldots, a_n, k \in \mathbb{Q}$; such formulas will often be abbreviated as $\sum_{\ell=1}^{n} a_\ell P_i(\varphi_\ell) \geq k$. Allowing for linear combinations of probability terms has a technical motivation (Fagin and Halpern, 1994). Usually, however, we will just be working with formulas like $P_i(\varphi) \geq k$, which should be read as: 'agent i assigns (subjective) probability (i.e. degree of belief) at least k to φ'. Furthermore, linear combinations of probabilities can be used

[8]The motivation for these two conditions is discussed extensively in Demey (2010, Section 3.1). The motivation for condition 1 is briefly hinted at later in this paper ('no dangerous divisions by 0').

to introduce conditional probabilities into the formal language. Recall that in probability theory, the conditional probability of φ given ψ is defined as follows:

$$P(\varphi|\psi) = \frac{P(\varphi \wedge \psi)}{P(\psi)} \quad \text{(provided } P(\psi) > 0).$$

It thus makes sense to introduce the following definition in the formal language:

$$P_i(\varphi|\psi) \geq k :\equiv P_i(\varphi \wedge \psi) - kP_i(\psi) \geq 0. \tag{4}$$

I will again abbreviate $[\![\varphi]\!]^{\mathbb{M}} := \{w \in W \mid \mathbb{M}, w \models \varphi\}$; furthermore, define $\mu_i(w)(X) := \sum_{x \in X} \mu_i(w)(x)$ for any $X \subseteq W$. The formal semantic clauses for proposition letters and the Boolean connectives are standard; the formal semantic clauses for the knowledge operator and probability formulas are the following:

$$\begin{array}{lll} \mathbb{M}, w \models K_i\varphi & \text{iff} & \forall v \in [w]_{\sim_i} : \mathbb{M}, v \models \varphi, \\ \mathbb{M}, w \models \sum_{\ell=1}^{n} a_\ell P_i(\varphi_\ell) \geq k & \text{iff} & \sum_{\ell=1}^{n} a_\ell \mu_i(w)([\![\varphi_\ell]\!]^{\mathbb{M}}) \geq k. \end{array}$$

A sound and complete axiomatization of this logic is readily available; cf. Fagin and Halpern (1994).

I now turn to the dynamic part. Given a probabilistic Kripke model $\mathbb{M} = \langle W, \sim_i, \mu_i, V \rangle_{i \in I}$ and a formula φ which is true in at least one state in $\mathbb{M}$, the *updated model* $\mathbb{M}|\varphi := \langle W^\varphi, \sim_i^\varphi, \mu_i^\varphi, V^\varphi \rangle$ is defined as follows:

- $W^\varphi = [\![\varphi]\!]^{\mathbb{M}}$,
- for every $i \in I$: $\sim_i^\varphi = \sim_i \cap (W^\varphi \times W^\varphi)$,
- for every $i \in I$ and $w \in W^\varphi$: $\mu_i^\varphi(w)(v) = \frac{\mu_i(w)(\{v\} \cap [\![\varphi]\!]^{\mathbb{M}})}{\mu_i(w)([\![\varphi]\!]^{\mathbb{M}})}$,
- for every $p \in Prop$: $V^\varphi(p) = V(p) \cap W^\varphi$.

(It is straightforward to check that the updated model $\mathbb{M}|\varphi$ is indeed a probabilistic Kripke model, as defined above.) Note that for the probabilistic part, the agents simply perform Bayesian updating (they conditionalize on the announced formula φ). Furthermore, note that μ_i^φ is always defined (no dangerous divisions by 0): $\mu_i^\varphi(w)$ is only defined for $w \in W^\varphi = [\![\varphi]\!]^{\mathbb{M}}$, so by condition 1 stated above it follows that $\mu_i(w)([\![\varphi]\!]^{\mathbb{M}}) \geq \mu_i(w)(w) > 0$. The language is extended with a dynamic operator $[!\,\cdot]$. Again, a formula such as $[!\varphi]\psi$ should be read as: 'after any public announcement of φ, it will be the case that ψ'. The dual of $[!\varphi]\psi$ is $\langle !\varphi\rangle\psi$, which should be read as: 'φ can be announced, and afterwards it will be the case that ψ'. The formal semantic clauses look as follows:

$$\begin{array}{lll} \mathbb{M}, w \models [!\varphi]\psi & \text{iff} & \text{if } \mathbb{M}, w \models \varphi \text{ then } \mathbb{M}|\varphi, w \models \psi, \\ \mathbb{M}, w \models \langle !\varphi\rangle\psi & \text{iff} & \mathbb{M}, w \models \varphi \text{ and } \mathbb{M}|\varphi, w \models \psi. \end{array}$$

Finally, one easily obtains a sound and complete axiomatization for this dynamified logic by adding reduction axioms. We focus on the reduction axioms for probability formulas. The reduction axiom for the formula $P_i(\psi) \geq \tau$ reads as follows:[9]

$$[!\varphi]P_i(\psi) \geq \tau \longleftrightarrow \Big(\varphi \to P_i(\langle !\varphi\rangle\psi \mid \varphi) \geq \tau\Big). \tag{5}$$

Note that to formulate a reduction axiom for the formula $P_i(\varphi) \geq \tau$, we used conditional probabilities. These can be defined in the formal language (cf. (4) above), so it is not strictly necessary to provide a separate reduction axiom for them.[10] Using (4) one easily obtains a reduction axiom for $P_i(\psi|\alpha) \geq \tau$:[11]

$$[!\varphi]P_i(\psi|\alpha) \geq \tau \longleftrightarrow \Big(\varphi \to P_i(\langle !\varphi\rangle\psi \mid \langle !\varphi\rangle\alpha) \geq \tau\Big). \tag{6}$$

[9]To achieve full generality, one needs to provide a reduction axiom not just for $P_i(\psi) \geq \tau$, but rather for $\sum_\ell a_\ell P_i(\psi_\ell) \geq \tau$ (involving linear combinations). This can easily be done; however, for our present purposes it will suffice to focus on the simpler case $P_i(\psi) \geq \tau$.

[10]In the previous subsection it *was* necessary to provide a separate reduction axiom for conditional belief, since *that* operator is *not* definable in the formal language.

[11]Again, it should be emphasized that (6) is an *axiom* in name only: it can be *derived* from the reduction axiom (5) for absolute probabilities and definition (4) of conditional probabilities in terms of absolute probabilities.

4.3 The Dynamics of the Lockean Thesis

Let's take stock. In Subsection 4.1 I discussed a system of public announcement logic, enriched with a qualitative notion of (conditional) belief. This system gives rise to the reduction axioms (2) and (3), for belief and conditional belief, respectively. In Subsection 4.2 I discussed probabilistic public announcement logic. This system gives rise to the reduction axioms (5) and (6), for probability and conditional probability, respectively. Note that (5) and (6) hold for any value of τ, so in particular when τ is 'high' (i.e. when $\tau > 0.5$).

Recall that the Lockean thesis says that belief can be defined as 'high' probability:

$$B_i\varphi :\equiv P_i(\varphi) \geq \tau$$

(where $\tau > 0.5$). A slightly more sophisticated version of this thesis says that *conditional* belief can be defined as 'high' *conditional* probability:

$$B_i(\varphi|\psi) :\equiv P_i(\varphi|\psi) \geq \tau.$$

If these two principles are applied to reduction axioms (5) and (6), then one simply obtains reduction axioms (2) and (3), respectively. In other words: if one accepts the Lockean thesis (and its slightly more sophisticated version), then the reduction axiom for high (conditional) probability is *exactly the same* syntactic expression as the reduction axiom for (conditional) belief. Accepting the Lockean thesis thus leads to a unified perspective on the dynamic behavior of belief and probabilities (degrees of belief).

What is the importance of this observation? One might be tempted to regard it as merely a technical 'artefact'. Nevertheless, it is quite surprising that the Lockean thesis leads to a unification between (2–3) and (5–6): it should be emphasized again that (2–3) are interpreted on epistemic plausibility models, which are purely *qualitative* entities (belief and conditional belief are interpreted by looking at $\leq_i$-minimal states; the definition of updated epistemic plausibility model is a straightforward extension of the well-known definition of updated Kripke model; etc.), whereas (5–6) are interpreted on probabilistic Kripke models, which have a large *quantitative* (probabilistic) component (belief and conditional belief, i.e. high probability and high conditional probability, are interpreted by means of the probability mass functions $\mu_i(w)$; the definition of updated probabilistic Kripke model essentially involves the idea of Bayesian conditionalization; etc.).

This seems to constitute a *pragmatic* or *methodological* argument in favor of the Lockean thesis. Accepting this thesis leads to a significant and unexpected unified perspective on the dynamic properties of technically very different frameworks. It thus helps to focus on the common purpose of these frameworks (despite their technical differences), viz. providing an account of 'soft information' and its dynamics. This is also relevant for practical or philosophical applications of these frameworks. For example, if in a given application one is heavily concerned with the dynamics of belief, but less so with its static properties (such as closure under conjunction), then both approaches described in this paper are equally applicable, and thus the final decision about which system to use will have to be motivated by other considerations.[12]

However, it might be possible to draw even further philosophical conclusions from our technical observation. Alexandru Baltag (2008; 2011) has argued for an 'Erlangen program' for epistemology: "in the spirit of Felix Klein's 1862 Erlangen program for mathematics, I argue that 'static' epistemic notions and properties are best characterized in terms of their transformations, their po-

[12]For example, Demey (2011a) analyzes the notion of surprise (focusing on its interaction with belief and its dynamic behavior) in a framework of probabilistic public announcement logic, with notions of belief and conditional belief defined according to the Lockean thesis. The reason for using the probabilistic framework (rather than the qualitative framework of epistemic plausibility models) is that probabilities are needed in the system for other, independent reasons as well, viz. as quantitative representations of *intensity of surprise*.

tential dynamics" (2011, p. 4). It was shown above that if one accepts the Lockean thesis (and its more sophisticated version)—if only for methodological reasons—, the epistemic notions of (conditional) belief and high (conditional) probability display exactly the same dynamic behavior (i.e. they have the same reduction axioms) with respect to public announcements. Baltag's Erlangen program for epistemology uses exactly this dynamic behavior to characterize epistemic notions, and therefore classifies (conditional) belief and high (conditional) probability as being one and the same epistemic notion. But this exactly means that the Lockean thesis should be accepted, not merely as a practically fruitful hypothesis, but also as a substantial epistemological claim about the notion of belief.

It might be objected at this point that belief and high probability really cannot be the same epistemic notion, simply because the former notion is closed under conjunction, whereas the latter isn't (cf. Section 3). However, from the perspective of Baltag's Erlangen program, this difference is a static difference (not a dynamic one), and should not be accepted as the sole criterium of individuation for epistemic concepts. With respect to dynamic behavior, which is deemed a more relevant individuation criterium, belief and high probability *do* have the same properties. In other words: the difference in closure under conjunction might indicate that belief and high probability are not the same notion *altogether*, but from an *epistemic* perspective they cannot be distinguished (the difference arises only at a *non-epistemic* level, for example the psychological level).

5 Conclusion

In this paper I have studied the Lockean thesis about beliefs and degrees of belief from the perspective of contemporary epistemic logic. The main problem of the Lockean thesis, viz. that it gives rise to a notion of belief which is not closed under conjunction, is typical for *classical* epistemic logic. I have argued that in *contemporary* epistemic logic, the Lockean thesis seems to have a much brighter future.

In the first place, I have briefly pointed out that this thesis can easily be extended from single-agent to multi-agent settings (via the notion of common p-belief). More importantly, however, I have shown that accepting it (and a more sophisticated version for conditional beliefs) leads to a significant and unexpected unification in the dynamic behavior of (conditional) belief (interpreted on epistemic plausibility models) and high (conditional) probability (interpreted on probabilistic Kripke models) with respect to public announcements. This already constitutes a strong argument for the methodological usefulness of the Lockean thesis. Furthermore, if one accepts Baltag's Erlangen program for epistemology, this technical observation has even stronger philosophical implications: because belief and high probability display the same dynamic behavior, it is very plausible that they are indeed one and the same epistemic notion.

Obviously, much more work needs to be done on this topic. In this paper it was shown that belief and high probability have the same dynamic behavior *with respect to public announcements*. However, for Baltag's Erlangen program to reach its full force, it is necessary to show that these two notions have the same dynamic behavior *in general*, i.e. with respect to an entire range of other types of dynamic phenomena. Secondly, there is a more philosophical issue that needs to be addressed. So far, Baltag's Erlangen program mainly seems to have a negative motivation: all attempts by classical epistemology to provide static definitions of the main epistemic notions (e.g. knowledge at time t is defined as justified true belief at time t)[13] have utterly failed, and therefore

[13] A notable exception is Goldman's (1979) 'historical reliabilism'. However, this theory is 'backward-looking' (epistemic states are characterized in terms of how they are generated),

it seems worthwile to look at an entirely new sort of individuation criterium, viz. sameness of dynamic behavior (this criterium has already proved to be successful in another area: geometry). Still, if this Erlangen program is to develop into a mature epistemological position, much more work will need to be done—in particular, providing a positive motivation.

Acknowledgements

I presented parts of this paper at Logic, Rationality and Reasoning (LRR) 10 (September 20–22, 2010, Ghent, Belgium) and at the First Philosophy Graduate Conference (April 1, 2011, Leuven, Belgium). I wish to thank the audiences of these talks for their stimulating questions and comments. I also wish to thank an anonymous referee for some useful suggestions. This research was financially supported by a PhD fellowship of the Research Foundation – Flanders (FWO).

References

Aumann, R. (1976). Agreeing to disagree. *Annals of Statistics*, 4:1236 – 1239.

Baltag, A. (2008). Interview. In Hendricks, V. F. and Pritchard, D., editors, *Epistemology: 5 Questions*, pages 21–37. Automatic Press.

Baltag, A. (2011). Research profile. Available online at the author's website: `http://alexandru.tiddlyspot.com/`.

Baltag, A. and Moss, L. S. (2004). Logics for epistemic programs. *Synthese*, 139:1–60.

whereas Baltag's proposal is 'forward-looking' (epistemic states are characterized in terms of how they can change).

Baltag, A. and Smets, S. (2008). A qualitative theory of dynamic interactive belief revision. In Bonanno, G., van der Hoek, W., and Woolridge, M., editors, *Texts in Logic and Games*, pages 9–58. Amsterdam University Press.

Barwise, J. (1988). Three views of common knowledge. In *TARK '88 Proceedings of the 2nd Conference on Theoretical Aspects of Reasoning about Knowledge*. Morgan Kaufmann.

Dégremont, C. and Roy, O. (2009). Agreement theorems in dynamic-epistemic logic. In He, X., Horty, J., and Pacuit, E., editors, *Logic, Rationality, and Interaction. LORI 2009 Proceedings*, LNAI 5834, pages 105 − 118. Springer.

Demey, L. (2010). Agreeing to disagree in probabilistic dynamic epistemic logic. Master's thesis, ILLC, Universiteit van Amsterdam.

Demey, L. (2011a). The dynamic phenomenon of surprise. An analysis in probabilistic dynamic epistemic logic. Manuscript.

Demey, L. (2011b). Some remarks on the model theory of epistemic plausibility models. *Journal of Applied Non-Classical Logics*, 21:375 − 395.

Douven, I. and Meijs, W. (2007). Measuring coherence. *Synthese*, 156:405–425.

Dubois, D. and Prade, H. (2009). Accepted beliefs, revision and bipolarity in the possibilistic framework. In Huber, F. and Schmidt-Petri, C., editors, *Degrees of Belief*, pages 161–184. Springer.

Eels, E. and Fitelson, B. (2000). Measuring confirmation and evidence. *Journal of Philosophy*, 97:663–672.

Fagin, R. and Halpern, J. (1994). Reasoning about knowledge and probability. *Journal of the ACM*, 41:340–367.

Foley, R. (1992). The epistemology of belief and the epistemology of degrees of belief. *American Philosophical Quarterly*, 29:111–121.

Goldman, A. (1979). What is justified belief? In Pappas, G. S., editor, *Justification and Knowledge*, pages 1–23. Dordrecht, Reidel.

Goldman, A. (1999). *Knowledge in a Social World*. Oxford University Press.

Halpern, J. Y. and Moses, Y. (1990). Knowledge and common knowledge in a distributed environment. *Journal of the ACM*, 37:549–587.

Halpern, J. Y., Samet, D., and Segev, E. (2009). Defining knowledge in terms of belief: the modal logic perspective. *Review of Symbolic Logic*, 2:469–487.

Hintikka, J. (1962). *Knowledge and Belief. An Introduction to the Logic of the Two Notions*. Cornell University Press.

Huber, F. and Schmidt-Petri, C., editors (2009). *Degrees of Belief*. Springer.

Jeffrey, R. (1983). *The Logic of Decision (2nd edition)*. University of Chicago Press.

Kajii, A. and Morris, S. (1997). Common p-belief: The general case. *Games and Economic Behavior*, 18:73–82.

Kooi, B. (2003). Probabilistic dynamic epistemic logic. *Journal of Logic, Language and Information*, 12:381–408.

Lewis, D. (1969). *Convention*. Harvard University Press.

Monderer, D. and Samet, D. (1989). Approximating common knowledge with common beliefs. *Games and Economic Behavior*, 1:170–190.

Spohn, W. (2009). A survey of ranking theory. In Huber, F. and Schmidt-Petri, C., editors, *Degrees of Belief*, pages 185–228. Springer.

van Benthem, J. (2007). Dynamic logic for belief revision. *Journal of Applied Non-Classical Logics*, 17:129–155.

van der Hoek, W., van Linder, B., and Meyer, J.-J. (1999). Group knowledge is not always distributed (neither is it always implicit). *Mathematical Social Sciences*, 38:215–240.

van Ditmarsch, H., van der Hoek, W., and Kooi, B. (2007). *Dynamic Epistemic Logic*. Springer.

Williams, M. (2001). *Problems of Knowledge. A Critical Introduction to Epistemology.* Oxford University Press.